arXiv:1610.06903v1 [cs.CV] 21 Oct 2016

# JOINT DEEP EXPLOITATION OF SEMANTIC KEYWORDS AND VISUAL FEATURES FOR MALICIOUS CROWD IMAGE CLASSIFICATION

*Joel Levis*[⋆]*, Hyungtae Lee*[†‡]*, Heesung Kwon*[‡]*, James Michaelis*[‡]*, Michael Kolodny*[‡]*, and Sungmin Eum*[‡§]

[⋆]Ohio University, Athens, Ohio, U.S.A.
[†]Booz Allen Hamilton Inc., McLean, Virginia U.S.A.
[‡]U.S. Army Research Laboratory, Adelphi, Maryland, U.S.A.
[§]University of Maryland, College Park, Maryland, U.S.A.

jl359113@ohio.edu, lee_hyungtae@bah.com, heesung.kwon.civ@mail.mil
james.r.michaelis2.civ@mail.mil, michael.a.kolodny.ctr@mail.mil, smeum@umiacs.umd.edu

## ABSTRACT

General image classification approaches differentiate classes using strong distinguishing features but some classes cannot be easily separated because they contain very similar visual features. To deal with this problem, we can use keywords relevant to a particular class. To implement this concept we have newly constructed a malicious crowd dataset which contains crowd images with two events, benign and malicious, which look similar yet involve opposite semantic events. We also created a set of five malicious event-relevant keywords such as *police* and *fire*. In the evaluation, integrating malicious event classification with recognition output of these keywords enhances the overall performance on the malicious crowd dataset.

***Index Terms***— malicious crowd dataset, semantic keyword, image classification

## 1. INTRODUCTION

General image classification methods have drawn upon the fact that images of differing classes have strong distinguishing features. [1, 2, 3, 4] However, certain classes involve very different events but can be represented with very similar image features, such as objects, that mainly appeared in the associated images. For example, in Figure 1, two images seem to contain similar event because persons are outstanding in both images. We can discern, however, that the two images involve opposite semantic events, which are benign and malicious. The right image is malicious due to several odd objects, such as smoke and police equipment. General image classification may not perform well without semantically crucial object information, which may or may not be notable from the image, but can use important keywords to guess which event occurs. We address this problem by identifying semantically unique keywords, which occur in higher frequency among the malicious images, and use these identified words to improve classification accuracy.

**Fig. 1**. A pair of similar looking crowd images with unique object contents

Since most benchmark datasets [5, 6, 7] collected for event classification do not deal with this problem, we collected a novel "malicious crowd" dataset, which contains crowd images with two events: benign and malicious. Along with event-level labels, we also collected a number of keywords that appeared in each image in the dataset, as listed below each image in Figure 1. We used Amazon Mechanical Turk to describe the semantic contents of each image in terms of keywords. Then we collected all the keywords for both classes and created a set of words used at most for each event. We selected non-overlapping distinctive keywords for the malicious event, which we aim to identify and treat them as the representative "semantic keywords".

To identify semantic keywords from a test image, we utilized a widely-used detection method, the deformable part model (DPM) [8], and a classification algorithm, which is a fine-tuned AlexNet [9]. Among the various keywords, words such as *police*, *helmet*, and *car* have rigid appearance but the others such as *fire* and *smoke* do not. DPM was used to detect the objects with rigid appearance whereas the finetuned AlexNet was employed to detect less rigid objects such as *smoke* and *fire*. We also built an additional fine-tuned AlexNet architecture to classify benign/malicious crowd images. Fi-

(a)

| category | keywords |
|---|---|
| benign | crowd, people, city, building, men, women, group, road, sidewalk, sign, race, tree, event, fans, gathering, . . . |
| malicious | crowd, people, protext, police, fire, street, riot, city, building, smoke, men, sign, flag, night, man, helmet, signs, group, violence, car, . . . |

(b)

**Fig. 2**. **Malicious Crowd Dataset:** (a) several example images for the benign and malicious events are shown in the first and second rows, respectively. (b) keywords mainly seen in the images for the benign and malicious images are listed. Red keywords are relevant keywords for the malicious event.

nally, we used several late fusion approaches to integrate the malicious crowd image classification result with the keyword detection/classification results. Our experiments show that fusion of image and keyword classifications outperforms the case when only the image classification is used. This supports the effectiveness of exploiting semantic keyword relevant to the malicious crowd images.

Our contributions are summarized as follows:

1. We introduce a new task of image classification where classes cannot be easily seperated from each other unlike general image classification.

2. To deal with this problem we collect a malicious crowd dataset which consists of two classes, malicious and benign crowds, which look similar but contain opposite semantic events.

3. We exploit semantic keywords only relevant to malicious crowd images to differentiate the malicious crowd images from the benign ones.

4. Integrating image features with these semantic keyword information increases image classification accuracy in the malicious crowd dataset.

## 2. MALICIOUS CROWD DATASET AND SEMANTIC KEYWORDS

### 2.1. Malicious Crowd Dataset

The "malicious crowd" dataset that was used to test our hypothesis contains 1133 crowd images equally split into two classes: benign and malicious. The intuition behind the labeling of the images was that, a benign crowd would be something a passerby would not be alarmed or concerned to see,

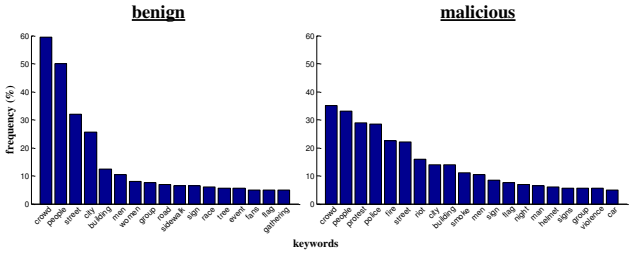


**Fig. 3**. **Histograms of relevant keywords:** The left and right histogram show the histograms of keywords relevant to benign and malicious classes, respectively. The keywords are listed according to their frequency of appearance in the images.

while a malicious image would be alarming and potentially dangerous.

These images were gathered using Google Images using various search terms. For benign images, search terms such as marathon, pedestrian crowd, parade, and concert were used. Riot and protest were used as search terms to gather the malicious crowd images. Figure 2(a) illustrates some example images from each class.

### 2.2. Semantic Keywords

To describe the contents of each of the crowd images, Amazon Mechanical Turk was used. A human was responsible for assigning five keywords to each image based on what objects are observed within. To ensure the accuracy of the Machanical Turk results, we manually removed the keywords which were incorrectly assigned.

After successfully collecting the crowd images and corresponding keywords, identifying keywords only relevant to the malicious class was necessary. We then constructed two keyword sets, each acquired by selecting the most frequently ap-

**Table 1**. Number of images where each keyword relevant to the malicious image appears.

| class | images | police | fire | smoke | helmet | car |
|---|---|---|---|---|---|---|
| benign | 557 | 8 | 1 | 2 | 7 | 57 |
| malicious | 576 | 205 | 144 | 150 | 206 | 65 |

pearing keywords in the two given classes. In practice, words that are commonly annotated in 5% or more images in each class were selected. As a result of this thresholding, the numbers of selected words for the benign and malicious classes are 17 and 20, respectively. Selected words and those frequency for both classes can be seen in Figure 3. We have refined the malicious keyword set by eliminating the keywords that appear in both classes. This elimination resulted in nine malicious keywords as shown in red in Figure 2(b). Lastly, we further eliminated keywords indicating particular phenomena such as *protest*, *riot*, *night*, and *violence*. Then *police*, *fire*, *smoke*, *helmet*, and *car* were included in the final set of malicious semantic keywords.

Table 2.2 shows the number of images where each keyword (object) actually appears. While *police*, *fire*, *smoke*, and *helmet* seem to be closely associated with the malicious event, *car* is seen in both events with a similar frequency. Note that the numbers in the table do not necessarily match the histogram of malicious semantic keywords obtained from Amazon Mechanical Turk. For example, *police* appears in 205 out of all 576 malicious images at a rate of 35.59%, but is assigned only to 28.50% of the malicious image by Amazon Mechanical Turk. This is because the visual contents associated with these keywords are not overly notable in several images. We can observe that the frequencies of the selected semantic keywords show a notable gap between the two classes, indicating that the purpose of the proposed keyword selection process is achieved.

## 3. THE PROPOSED APPROACH

To identify semantic keywords from the test images, keyword detectors/classifiers were trained. For objects with rigid appearance such as *police*, *helmet*, and *car*, deformable part models (DPM) [8] were trained. For *fire* and *smoke* which are objects with non-rigid appearance, convolutional neural network (CNN) classifiers finetuned on the AlexNet architecture [9] were used. Since the object detectors output multiple detections for an image, we select one detection with a maximum score and use that score to represent the confidence of the object presence in that image. We also built a CNN classifier to output the confidence score for the maliciousness of an image. Multiple late fusion approaches were utilized to combine the output of all keyword detectors/classifiers and the malicious image classifier.

### 3.1. Learning Keyword Detectors

DPM detectors which are used to identify *police* and *helmet* were trained on 400 annotated images, made up of all auxiliary images from Google Images. For a *car* detector, we used the DPM trained on PASCAL VOC 2007 dataset [10].

### 3.2. Learning Malicious Event/Keyword Classifiers

Firstly, a finetuned AlexNet deep convolutional neural network (DCNN) was trained to classify images as benign or malicious. The training set includes 905 images randomly selected from the malicious crowd dataset. Finetuning was conducted on all eight layers of AlexNet, with the eighth layer learning with a learning rate of 20 and a learning rate of 2 for all others. The last layer was replaced so as to have a binary output in contrast to the 1000 class output of AlexNet.

The *fire* and *smoke* DCNN-based classifiers were also trained in a similar way to the previously described DCNN. Each of these models was trained on 300 images. These contain images from our dataset and the auxiliary images gathered from Google Images. We used seperate networks for the two keywords instead of one network with multiple labels because both keywords may appear in the same training image.

### 3.3. Late Fusion

A late fusion was performed on the output of the six streams which are the malicious crowd image classifier, three detectors for *police*, *helmet*, and *car*, and two classifiers for *fire* and *smoke*. The late fusion is used to enhance the baseline classifier with the thought that additional object information would help to increase classification accuracy. In an attempt to test which fusion method would be most effective, the streams were tested using various fusion methods. These included Linear Discriminant Analysis (LD) [11], Logistic Regression (LR) [12], Support Vector Machines (SVM) [13], $k$-Nearest Neighbor Classifiers ($k$NN) [14], Subspace-based Ensemble Classifiers (EC) [15], and a Dynamic Belief Fusion (DBF) [16]. For SVM, we used two different kernels which are a linear kernel (SVM-lin) and RBF kernel (SVM-rbf). For $k$NN, we used 100 clusters and these clusters are clustered according to the Euclidean distance. As the EC, we used a subspace ensemble classifier with a set of 30 weak models.

## 4. EXPERIMENTS

### 4.1. Dataset Partition and Evaluation Protocol

The Malicious Crowd Dataset consists of 1133 images - 576 of 1133 are labeled as the malicious crowd image and the rest are labeled as benign. The same training dataset mentioned in Section 3.2 (905 images) is used to train the fusion approaches. The rest (228 images) are used as the test set.

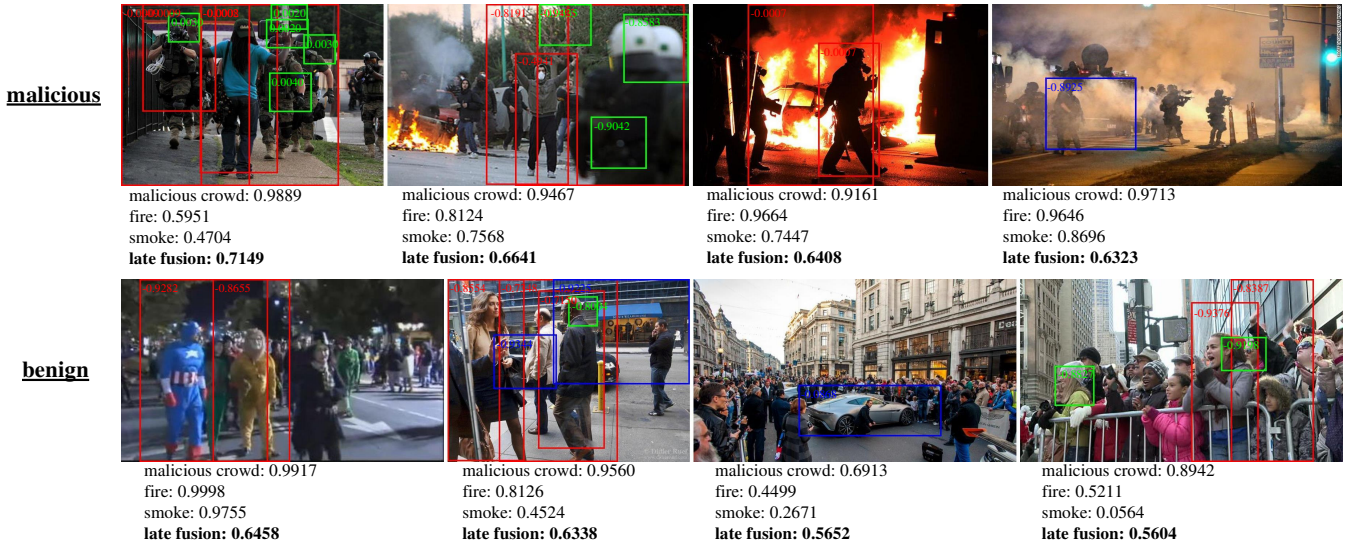


**Fig. 4**. **Output of malicious crowd image classification and keyword detectors/classifiers:** the first and second row show four examples with largest fusion scores for malicious event from malicious and benign crowd images, respectively. Bounding box with color of red, green, and blue indicates detection of *police*, *helmet*, and *car* detector, respectively. A late fusion score is obtained by EC (a subspace ensemble classifier).

**Table 2**. Malicious crowd image classification accuracy measured by AP

| | | keyword | | | | | late fusion | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | baseline | police | fire | smoke | helmet | car | SVM-rbf | DBF | SVM-lin | kNN | LD | LR | EC |
| AP | .722 | .586 | .563 | .689 | .532 | .491 | 742 | .757 | .758 | .758 | .760 | .763 | **.771** |
| Gain | · | · | · | · | · | · | +.020 | +.035 | +.036 | +.036 | +.038 | +.041 | +.049 |

Averge precision (AP) is used as an evaluation metric in our experiments.

### 4.2. Results

Table 2 shows a malicious crowd image classification accuracy in AP for a baseline malicious crowd image classification, keyword detections/classifications, and various late fusion approaches. Note that, for a keyword detection/classification, classification accuracy was calculated for recognizing malicious image instead of each associated keyword. For example, when the test image is originally malicious while not containing any police in the image, if the *police* detector does not detect any police in the image, the result is still considered false negative. Using the *car* detector does not provide competitive accuracy because, as shown in the Table 2.2, *car* is not significantly relevant to the malicious crowd. Other keyword detectors do not provide better classification accuracy than baseline malicious crowd image classification. This is because these sematic keywords (objects) are only seen in small portions in the dataset. However, integrating the baseline with the output of these keyword classifiers/detectors enhanced the classification accuracy by approximately 7% at most. The best performer is EC, a subspace-based ensemble classifier, achieving fusion gain of .049 in AP. We can observe that all fusion approaches improve classification accuracy over the baseline, which supports the benefit of jointly exploiting semantic keywords and the associated detectors and classifiers. Figure 4 shows several images with high scores in terms of their maliciousness from both malicious and benign classes.

## 5. CONCLUSIONS

We addressed the new image classification problem where certain classes can be expressed by similar visual features but should be distinguished from each other semantically. To demonstrate, we have constructed a novel malicious crowd image dataset which consists of two classes (benign and malicious) that may look similar but contain semantically different events. To better classify the images with the aforementioned characteristics, we have selected representative keywords for malicious crowd images which are then incorporated with conventional image classifiiers using a multi-stream late fusion architecture. As Table 2 shows, the approach that we have hypothesized lead to considerable performance improvements over the conventional baseline classifier when used in practice.

## 6. REFERENCES


[1] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton, "Imagenet classification with deep convolutional neural networks," *Advances in Neural Information Processing Systems 25*, 2012.

[2] Maxime Oquab, Léon Bottou, Ivan Laptev, and Josef Sivic, "Is object localization for free? – weakly-supervised learning with convolutional neural networks," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015.

[3] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.

[4] Archith J. Bency, Heesung Kwon, Hyungtae Lee, S Karthikeyan, and B. S. Manjunath, "Weakly supervised localization using deep feature maps," *European Conference on Computer Vision*, 2016.

[5] Li-Jia Li and Li Fei-Fei, "What, where and who? classifying event by scene and object recognition," *IEEE International Conference on Computer Vision*, 2007.

[6] Sangmin Oh, Anthony Hoogs, Amitha Perera, Naresh Cuntoor, Chia-Chih Chen, Jong Taek Lee, Saurajit Mukherjee, JK Aggarwal, Hyungtae Lee, Larry Davis, et al., "A large-scale benchmark dataset for event recognition in surveillance video," *IEEE Conference on Computer Vision and Pattern Recognition*, 2011.

[7] George Awad, Jonathan Fiscus, Martial Michel, David Joy, Wessel Kraaij, Alan F. Smeaton, Georges Quéenot, Maria Eskevich, Robin Aly, and Roeland Ordelman, "Trecvid 2016: Evaluating video search, video event detection, localization, and hyperlinking," in *Proceedings of TRECVID 2016*. NIST, USA, 2016.

[8] Pedro F. Felzenszwalb, Ross B. Girshick, David McAllester, and Deva Ramanan, "Object detection with discriminatively trained part based models," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 9, pp. 1627–1645, 2010.

[9] Maxime Oquab, Léon Bottou, Ivan Laptev, and Josef Sivic, "Learning and transferring mid-level image representations using convolutional neural networks," *IEEE Conference on Computer Vision and Pattern Recognition*, 2014.

[10] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John Winn, and Andrew Zisserman, "The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results," http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html.

[11] Ronald Alymer Fisher, "The use of multiple measurements in taxonomic problems," *Annals of Eugenics*, vol. 7, pp. 179–188, 1936.

[12] David A. Freedman, "Statistical models: Theory and practice," p. 128. Cambridge University Press, 2009.

[13] C Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.

[14] N. S. Altman, "An introduction to kernel and nearest-neighbor nonparametric regression," *The American Statistician*, vol. 46, no. 3, pp. 175–185, 1992.

[15] Tin Kam Ho, "The random subspace method for constructing decision forests," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 20, no. 8, pp. 832–844, 1998.

[16] Hyungtae Lee, Heesung Kwon, Ryan M. Robinson, William D. Nothwang, and Amar M. Marathe, "Dynamic belief fusion for object detection," *IEEE Winter Conference on Applications of Computer Vision*, 2016.